arXiv:1408.3783v2 [cs.GT] 18 Sep 2014

Long-term causal effects of economic mechanisms on agent incentives

Panos Toulis and David C. Parkes

September 19, 2014

Abstract

Economic *mechanisms* administer the allocation of resources to interested agents based on their self-reported *types*. One objective in *mechanism design* is to design a *strategyproof* process so that no agent will have an incentive to misreport its type. However, typical analyses of the incentives properties of mechanisms operate under strong, usually untestable assumptions. Empirical, data-oriented approaches are, at best, under-developed. Furthermore, mechanism/policy evaluation methods usually ignore the dynamic nature of a multi-agent system and are thus inappropriate for estimating long-term effects. We introduce the problem of estimating the *causal effects* of mechanisms on incentives and frame it under the Rubin causal framework (Rubin, 1974, 1978). This raises unique technical challenges since the outcome of interest (agent truthfulness) is confounded with strategic interactions and, interestingly, is typically never observed under any mechanism. We develop a methodology to estimate such causal effects that using a prior that is based on a strategic equilibrium model. Working on the domain of *kidney exchanges*, we show how to apply our methodology to estimate causal effects of kidney allocation mechanisms on hospitals' incentives. Our results demonstrate that the use of game-theoretic prior captures the dynamic nature of the kidney exchange multi-agent system and shrinks the estimates towards long-term effects, thus improving upon typical methods that completely ignore agents' strategic behavior.

Keywords: causal inference, multiagent systems, equilibrium effects, mechanism design

1 Introduction

A mechanism defines the rules that are used to determine the allocation of resources to interested parties in an economic transaction. For example, an online ad auction determines the winners of the advertisement slots and the appropriate payments, based on advertisers' reports (bids). In designing mechanisms, a key requirement is to have good incentives properties, so that agents have no incentive to misreport their valuations. Such strategyproof mechanisms are appealing for being strategically simple for agents, and can lead to desirable outcomes because this agent behavior can be anticipated and leveraged for good effect.

There are a few general procedures for devising strategyproof mechanisms. One key idea is to determine the payment of an agent i as a function of the reports of all agents excluding i. Along with allocating the desired resources to this agent given the prices it faces, the intuition is that agent i will have no incentive to misreport since this will not affect the utility given fixed reports from others. This idea underlies the Vickrey auction (Vickrey, 1961) and its generalization in the Vickrey-Clarke-Groves (VCG) mechanism (Clarke, 1971; Groves, 1973). For example, in the Vickrey auction, the highest bidder wins the item but pays the second highest bid and thus, there is no incentive for the highest bidder to reduce the initial bid.

However, and even when strategyproof mechanisms can in principle be designed, theoretical analyses of their incentive properties rely critically on assumptions that are strong or untestable in practice. Typical assumptions include: no *collusion* among agents; (ii) the *rationality* of participants; (iii) that the types (such as valuations) of agents are correctly modeled (e.g., they are private values and don't depend on information of others); and (iv) that the strategic interactions have been correctly modeled. But without getting these assumptions correct, the incentive properties of a mechanism will not be as desired. For example, if participants in a single-item second price (Vickrey) auction can collude then one bidder can submit a high bid while others withhold their bids. In another example, if the problem is truly multi-round then participants have a new incentive to shave down their bids in order to get the best price for an item across time.

In this light, there remains a large opportunity for empirical methods in the design of mechanisms, especially in estimating the *causal effects* of mechanisms on incentives, and other outcomes of interest (such as welfare, revenue and so forth.) Across many online platforms such as ad auction platforms, one wants to be able to make changes to the design across subsets of the population and be able to estimate the effect of these design decisions on economic properties if one was to run a single design on the whole population vs run some other design. In this paper we adopt the Rubin causal framework (Rubin, 1974, 1978) using the potential outcomes notation. Our goal is to estimate the causal effects of mechanisms on the incentives of agents, after agents have

been randomly assigned to participate in one mechanism (viewed as a treatment) and their reports have been observed. This raises unique technical challenges since the outcome of interest (agent truthfulness) is typically never observed under any mechanism. Furthermore, we assume that data collection (agent reports) happens before the system has reached an equilbrium and so, we are interested in a methodology that will strike a sensible balance between observed data and equilbrium considerations.

There is a developing body of work on experimention of online, socio-economic systems, but it is not developed within the potential outcomes framework and has not studied the special question of causal analysis in regard to incentive properties. In related work, large field experiment conducted at Yahoo! in 2008, aimed to estimate the effects of increased reserve prices on keyword revenue (Ostrovsky & Schwarz, 2011). The applied method was to use a "diff-in-diffs" estimator that completely ignores all aforementioned subtleties. Other work aims to estimate the effects of interventions in a machine learning model underlying a mechanism (e.g. (Bottou et al., 2012)), but the methods are usually predictive (i.e. predict all missing outcomes through a model based on one intervention and the observed outcomes). Equilibrium effects for causal inference has first been proposed in the econometric literature (Heckman & Vytlacil, 2005). However, no general methodology has been proposed for the estimation of long-term causal effects of policies/mechanisms.

Other work has considered the empirical design of mechanisms in settings where the goal of strategyproof design is unachievable in combination with other desirable properties, or cannot be supported from an analytical framework (Lubin & Parkes, 2009). This appeals to the divergence between distributions over payments (or payoffs) in an incentive-aligned mechanism and distributions over payments (or payoffs) in another candidate design, with the view to finding the optimal mechanism through online search. But this work does not adopt a causal approach, but rather assumes the ability to switch the entire population through alternate designs and thus does not have the difficulty of estimating counterfactuals. Moreover, the work does not adopt our viewpoint of looking to make inferences from empirical frequences about reported types about the incentive properties of the mechanism.

2 Causal effects on incentives

We consider a population of N agents, indexed by i in some natural ordering, and two mechanisms $\mathcal{M}_1$ and $\mathcal{M}_0$. Each agent will be randomized to participate in one mechanism only (thus the mechanisms can be viewed as treatments that agents receive). Specifically, if agent i participates in $\mathcal{M}_1$ then $Z_i = 1$ and if agent i participates in $\mathcal{M}_0$ then $Z_i = 0$. We also consider a setting

where the mechanisms are multi-round, each round being indexed by t, for $t = 1, 2, \cdots T$ and a maximum number of rounds T. At each round t of any mechanism, each agent samples a type θ_{it} from a type space Θ according to some distribution F and selects a strategy $Y_i(z,t) \in \{0,1\}$. Given the sampled type, agent i, then reports $R_i(z,t)$ to mechanism $\mathcal{M}_z$ according to the following rule:

$$R_i(z,t) = \begin{cases} \theta_{it}, & \text{if } Y_i(z,t) = 1 \\ d(\theta_{it}) \in \Theta, & \text{otherwise} \end{cases}$$

In other words, if $Y_i(z,t) = 1$, agent i is reporting truthfully to the mechanism and it is deviating according to a known deviation if $Y_i(z,t) = 0$. We assume that, the distribution of true types F and the function of deviation $d(\cdot)$, are known or can be estimated from other sources[1]. Thus, assume that $d(\theta_{it})$ has a known distribution G. The agent strategy $Y_i(z,t)$ and report $R_i(z,t)$ at each round t are the potential outcomes of interest.

We consider a completely randomized experiment and denote the entire $(N \times 1)$ assignment vector with $\mathbf{Z}$. The full $(N \times 1)$ vector of potential outcomes of agent stratagies for mechanisms $\mathcal{M}_0$ and $\mathcal{M}_1$ at round t, are denoted by $\mathbf{Y}_0(t)$ and $\mathbf{Y}_1(t)$ respectively. In a similar fashion, let $\mathbf{R}_0(t), \mathbf{R}_1(t)$ denote the potential reports of agents in mechanisms $\mathcal{M}_0, \mathcal{M}_1$ respectively. Note that all the potential outcomes for agent strategies, $\mathbf{Y}_0(t), \mathbf{Y}_1(t)$, are never observed and thus are considered as missing. However, we make the following distinction: for an agent i with $Z_i = z \in \{0,1\}$, the outcome $Y_i(z,t)$ will be *realized* but will not be observed whereas the outcome $Y_i(1-z,t)$ will not be realized at all. The subvector of $\mathbf{Y}_0(t)$ with the realized outcomes for $\mathcal{M}_0$ is denoted by $\mathbf{Y}_0^{\text{real}}(t)$. Similarly, the vector of realized outcomes for $\mathcal{M}_1$ is denoted by $\mathbf{Y}_1^{\text{real}}(t)$. In contrast, some of the potential outcomes of the reports of agents are observed under the mechanisms they participate in. Let $\mathbf{R}_0^{obs}(t)$ denote the subvector of $\mathbf{R}_0(t)$, for those agents i such that $Z_i = 0$, and $\mathbf{R}_1^{obs}(t)$ be the subvector of $\mathbf{R}_1(t)$ for agents i with $Z_i = 1$.

The "Science" table of the observed and unobserved quantities in the aforementioned experiment is shown in Table 1. We can now define the causal estimand of interest:

Definition 2.1 (Causal effects on incentives)**.** *The causal effect on incentives of mechanism $\mathcal{M}_1$ over mechanism $\mathcal{M}_0$ in round t, is defined by:*

$$\Delta(t) = \frac{1}{N} \sum_{i=1}^{N} Y_i(1,t) - Y_i(0,t) \tag{1}$$

[1]For example, in the ad auction literature bidder valuations can be routinely estimated from the data using empirical distribution of bids and prices (Athey & Nekipelov, 2010; Ostrovsky & Schwarz, 2011)

Table 1: Science table of potential outcomes. Outcomes of agent strategies $Y_i(z,t)$ are all missing and reports $R_i(z,t)$ are only observed when $Z_i = z$.

		$\mathcal{M}_0$		$\mathcal{M}_1$	
Units	$\mathbf{Z}$	$\mathbf{R}_0(t)$	$\mathbf{Y}_0(t)$	$\mathbf{R}_1(t)$	$\mathbf{Y}_1(t)$
1	0	$R_1(0,t)$	?	?	?
2	0	$R_2(0,t)$	?	?	?
...		...		...	
N-1	1	?	?	$R_{N-1}(1,t)$	?
N	1	?	?	$R_N(1,t)$	?

The estimand $\Delta(T)$ is defined as the long-term effect on agent incentives.

2.1 Discussion

Recall that $Y_i(z,t)$ denotes the strategy of agent i (1=truthful, 0=deviating) it mechanism $\mathcal{M}_z$ at round t. Therefore, the estimand $\Delta(t)$ defined in (1) compares the proportions of truthful agents in $\mathcal{M}_1$ compared to $\mathcal{M}_0$ and by definition, it holds that $\Delta(t) \in [-1, 1]$. Other options for the definition of the estimand are available as well (e.g. median of difference) and in general it would involve a "contrast function" $h(\mathbf{Y}_1(t), \mathbf{Y}_0(t))$ that will summarize the difference between the two vectors. We will use this notion of contrast function throughout the rest of this paper, but for all numerical purposes we will assume this is the difference in means as in Definition 1.

Note also that the estimand is time-dependent as we expect agents to be self-interested and adapt their strategy over repeated rounds. The long-term effect $\Delta(T)$ is trying to capture this dynamic evolution of agent strategies over a specified time horizon T that is considered enough time for the system to reach an economic equilbrium. This is related to the study of *equilibrium effects* in the econometric literature (Heckman & Vytlacil, 2005).

3 Causal Inference

Inspection of Table 1 reveals one major technical challenge. Since we do not observe the actual strategies of agents (i.e., their "truthfulness" status) but only their reports, no potential outcomes of strategies $Y_i(z,t)$ are actually observed. However, given strategies $Y_i(z,t)$ the potential outcomes

on reports $R_i(z,t)$ have a well-defined distribution[2].

$$p(\mathbf{Y}_1(t), \mathbf{Y}_0(t)|\mathbf{R}_1^{obs}(t), \mathbf{R}_0^{obs}(t), \mathbf{Z}) \propto \mathcal{L}(\mathbf{R}_1^{obs}(t), \mathbf{R}_0^{obs}(t)|\mathbf{Y}_0(t), \mathbf{Y}_1(t)) \times \pi(\mathbf{Y}_0(t), \mathbf{Y}_1(t)) \tag{2}$$

The likelihood term $\mathcal{L}(\mathbf{R}_1^{obs}(t), \mathbf{R}_0^{obs}(t)|\mathbf{Y}_0(t), \mathbf{Y}_1(t))$ is easy to obtain based on our assumptions. Specifically, we have assumed that report $R_i(z,t)$ has distribution F if agent i is truthful and has distribution G if it is deviating. Hence:

$$\mathcal{L}(\mathbf{R}_1^{obs}(t), \mathbf{R}_0^{obs}(t)|\mathbf{Y}_0(t), \mathbf{Y}_1(t)) = \mathcal{L}(\mathbf{R}_1^{obs}(t)|\mathbf{Y}_1(t)) \times \mathcal{L}(\mathbf{R}_0^{obs}(t)|\mathbf{Y}_0(t))$$

Therefore, by independence, it holds for $j \in \{0,1\}$ indexing mechanism $\mathcal{M}_j$:

$$\mathcal{L}(\mathbf{R}_j^{\text{obs}}(t)|\mathbf{Y}_j^{\text{real}}(t)) = \prod_{i:Z_i=j} f(R_i(j,t))^{Y_i(j,t)} g(R_i(j,t))^{1-Y_i(j,t)} \tag{3}$$

Hence, causal inference depends critically on the model of $\pi(\mathbf{Y}_0(t), \mathbf{Y}_1(t))$. The main contribution of this paper is to consider a prior on potential outcomes that has a game-theoretic justification through a well-defined equilibrium model. The main idea is that, by doing so, we will shrink estimates from data observed at an early round towards the long-term effects, assuming that the equilibrium model is accurate enough to describe the dynamics of the economic system. To illustrate our method we will compare it with a straightforward imputation method the is based on a uniform prior. More options, such as a fully-Bayesian approach are discussed later.

3.1 Empirical method: Imputation on uniform prior of realized outcomes

This method, dubbed the *empirical method*, serves as our baseline method and works in two steps. First, we impute the realized (but missing outcomes) $\mathbf{Y}_1^{\text{real}}(t), \mathbf{Y}_0^{\text{real}}(t)$ assuming a uniform prior. Second, we impute the non-realized and missing outcomes $\mathbf{Y}_0^{\text{nreal}}(t), \mathbf{Y}_1^{\text{nreal}}(t)$ through the empirical distribution of the imputed realized outcomes. This algorithmic process (shown next) is repeated many times and estimates of the causal effects are used for summarization.

[2]Since we operate under a completely randomized experiment, the assignment mechanism $p(\mathbf{Z}|\cdots)$ is unconfounded and the vector $\mathbf{Z}$ can be omitted for brevity.

Initialize Δ = array of length n
For $j = 1, 2, \cdots n$
    Impute all missing potential outcomes for strategies as follows:
        $\mathbf{Y}_1^{\text{real}}(t) \sim \mathcal{L}\left(\mathbf{R}^{\text{obs}}(1,t)|\mathbf{Y}_1^{\text{real}}(t)\right)$
        $\mathbf{Y}_1^{\text{nreal}}(t) \sim \mathbf{Y}_1^{\text{real}}(t)$ (use empirical distribution)
        $\mathbf{Y}_0^{\text{real}}(t) \sim \mathcal{L}\left(\mathbf{R}^{\text{obs}}(0,t)|\mathbf{Y}_0^{\text{real}}(t)\right)$
        $\mathbf{Y}_0^{\text{nreal}}(t) \sim \mathbf{Y}_0^{\text{real}}(t)$ (use empirical distribution)
    $\mathbf{Y}_1(t) = (\mathbf{Y}_1^{\text{nreal}}(t), \mathbf{Y}_1^{\text{real}}(t))$
    $\mathbf{Y}_0(t) = (\mathbf{Y}_0^{\text{real}}(t), \mathbf{Y}_0^{\text{nreal}}(t))$
    Causal effect estimate $\hat{\Delta}_j(t) = h(\mathbf{Y}_1(t), \mathbf{Y}_0(t))$
Return Δ

Algorithm 1: Causal inference on incentives through uniform priors. Unrealized outcomes are sampled from the empirical distribution of the imputed realized ones. Variable n is the # of samples.

The estimate of causal effects on incentives from the empirical method is given by:

$$\widehat{\Delta(t)} = \frac{1}{n}\sum_j \hat{\Delta}_j(t) \tag{4}$$

One critical implicit assumption underlying the imputation of non-realized outcomes from realized ones, is that collective behavior is somehow "homegeneous" in a mechanism. For example, if 2 out of 10 agents are truthful on average, then we expect 4 agents to be truthful out of 20.

3.2 Game-theoretic method: Imputation using a game-theoretic prior

Typical causal inference methods, such as the aforementioned one, are usually criticized that they ignore incentives in a multi-agent system (Heckman & Vytlacil, 2005). However, in this work we show that the Rubin causal model can be adapted to address this issue. Before we proceed, we make the following assumption:

Assumption 3.1 (Best-response behavior)**.** *Given no prior information, the potential outcomes on agent strategies are independent for every round i.e.,* $\mathbf{Y}_1(t) \perp\!\!\!\perp \mathbf{Y}_0(t), \forall t$.

Assumption 3.1 can be thought as a consequence of assuming that agents are best-responding, regardless of behaviors in other mechanisms. This assumption would be invalid in several cases, for example, when agents have different propensities to be truthful or lie to a mechanism. We will offer more discussion later in the paper.

Assuming independence, we only need to model $\pi(\mathbf{Y}_0(t))$ and, similarly, $\pi(\mathbf{Y}_1(t))$. Recall that in any mechanism, say $\mathcal{M}_1$, the expected utility of agent i choosing strategy $Y_i(z,t) = y$ assuming fixed behaviors from other agents $\mathbf{Y}_{1,-i}(t)$ is denoted by $u_i(y, \mathbf{Y}_{1,-i}(t))$. Therefore the expected utility benefit from being truthful for agent i is given by

$$\Delta u_i(\mathbf{Y}_{1,-i}(t)) = u_i(1, \mathbf{Y}_{1,-i}(t)) - u_i(0, \mathbf{Y}_{1,-i}(t)) \tag{5}$$

We adopt a *quantal response equilbrium* (McKelvey & Palfrey, 1995; Goeree et al., 2003) model in order to construct our game-theoretic prior. In specific, agent i facing a vector of agent strategies, will randomize over the available actions according to a softmax rule that depends on the expected utilities. In specific, for a mechanism $\mathcal{M}_z$ at round t, agent i facing fixed behaviors from other agents $\mathbf{Y}_{z,-i}(t)$, will select to be truthful according to the probability:

$$\pi(Y_i(z,t) = 1 | \mathbf{Y}_{z,-i}(t)) \propto \exp\left(\beta \cdot u_i(1, \mathbf{Y}_{z,-i}(t))\right) \tag{6}$$

Hence:

$$\pi(Y_i(z,t) = 1 | \mathbf{Y}_{z,-i}(t)) = \frac{1}{1 + \exp\left(\beta \cdot \Delta u_i(\mathbf{Y}_{z,-i}(t))\right)}$$

Quantal response equilbrium is a well-studied model of utility-based agent behavior that has been shown to converge to Nash equilibria under certain mild conditions (McKelvey & Palfrey, 1995). The choice of parameter $\beta > 0$ is critical[3]. If β is high then the agent has a strong preference for actions with better expected utilities. In the extreme case, if β is very high then the agent simply prefers the best action (so adopts a best-response strategy). We will discuss the choice of β in the experimental section. Note also that there are N functions $u_i(\cdot)$ and each needs to be evaluated at 2^{N-1} points since the outcome is binary. For large populations this computation is prohibitive. To circumvent this problem we need to make the following exchangeability assumption:

Assumption 3.2 (Exchangeability among agents)**.** *Agents are exchangeable so that inferences are invariant to permutations of agent labels.*

The main consequence of the exchangeability assumption is that the expected utility $u_i(\cdot)$ are the same for all agents. Furthermore, the expected utility of an agent depends only on the number of truthful agents he is "competing" against and so $\Delta u_i(\mathbf{Y}_{z,-i}(t)) = \Delta u(\sum_{j \neq i} Y_j(z,t))$. In other words, there is only one function $\Delta u(\cdot)$ for each agent that needs to be evaluated at $N + 1$

[3]If $\beta < 0$ that would be considered irrational since the agent would prefer actions with smaller expected utilities than others.

points (having 0 to N truthful agents in total). Inference can now be performed through Gibbs sampling. The implementation is straightforward since strategy outcomes are binary. In summary, if i participates in mechanism $\mathcal{M}_z$, then potential outcome $Y_i(z,t)$ is sampled according to the following rule: If the outcome is not-realized then only the report outcome $R_i(z,t)$ is used (this is observed) to sample the strategy. However, if the outcome is realized, then both the report $R_i(z,t)$ and the vector of agent behaviors $\mathbf{Y}_{z,-i}(t)$ is used to sample $Y_i(z,t)$. The full procedure is given in Algorithm 2.

Initialize $\Delta =$ array of length n
For $j = 1, 2, \cdots n$
  Initialize $\mathbf{Y}_0^{(j)}(t) = \mathbf{Y}_0^{(j-1)}(t)$ and $\mathbf{Y}_1^{(j)}(t) = \mathbf{Y}_1^{(j-1)}(t)$
  $\forall i \in \{1, 2, \cdots, N\}$ and $z \in \{0, 1\}$:

$$\pi_{iz} = \left(1 + \exp\left(\beta \Delta u(\textstyle\sum_{j \neq i} Y_j(z,t))\right)\right)^{-1}$$

$$p_{iz} = \begin{cases} \mathcal{L}(R_i(z,t) | Y_i^{(j)}(z,t) = 1), & \text{if } Z_i = 1 - z \\ \mathcal{L}(R_i(z,t) | Y_i(z,t) = 1) \times \pi_{iz}, & \text{if } Z_i = z \end{cases}$$

  Sample $Y_i^{(j)}(z,t) = 1$ with probability p_{iz}
  Causal effect estimate $\hat{\Delta}_j(t) = h(\mathbf{Y}_1^{(j)}(t), \mathbf{Y}_0^{(j)}(t))$
Return Δ

Algorithm 2: Causal inference on incentives through a game-theoretic prior. Inference performed though Gibbs sampler. Each sample is indexed by j. Variable n is the # of samples.

4 Application on Kidney exchanges

4.1 Preliminaries

Kidney exchanges (Roth et al., 2004) enable kidney transplantations when donors are incompatible with recipients. In particular, a pair of a donor and a recipient who are incompatible can exchange a kidney transplant with a pair of donor/recipient, provided that the donor from one pair can donate to the patient of the other. Incompatilibility is determined by two medical tests. The first is one is a blood-type test between the donor and the recipient. The second test is a sensitivity test which shows whether the recipient will accept or reject the kidney transplant from the donor. The statistics of these compatibilities are well studied and we will assume them to be known [4]. Typically, these exchanges involve 2 pairs due to logistical issues, however it is also common to perform *cycle*

[4]For example, it is known that the probability that a random patient will reject a kidney of a random donor is about 0.11 and this is 3x as high when the recipient is a woman who has been pregnant and the donor is her spouse.

exchanges in which a donor donates to the recipient of another pair, in sequence, until a loop is formed. Multiple regional exchange programs currently operate in the US and the world, however, their expansion has hitherto been hindered by logistical and mechanism inefficiency issues. In specific, it has been reported that manipulation of centralized kidney exchange markets is possible and is performed by participating hospitals (Ashlagi et al., 2010). Work in mechanism design has focused on mechanisms that resolve such incentives issues (Ashlagi & Roth, 2011; Toulis & Parkes, 2011; Ashlagi et al., 2010).

The kidney exchange problem fits the framework of this paper as follows. First, we assume N hospitals and we assume the existence of two mechanisms $\mathcal{M}_0$ and $\mathcal{M}_1$. The former, $\mathcal{M}_0$, can be considered as the mechanism currently in practice whereas $\mathcal{M}_1$ is a new proposed mechanism under test. Agents are hospitals that are randomly assigned to participate in an exchange mechanism. This exchange is multi-round (e.g. once per month) as it usually happens in practice. At each round t, each hospital *samples* a donor/patient pool of fixed size. This pool can be represented by a set of donor/patient pairs such that in each pair the donor is willing to donate to the patient but they are incompatible. Given such set, compatibilities can be determined by medical tests that are assumed *common knowledge* i.e., hospitals cannot hide compatibilities between pairs, as that would be unethical and easy to uncover. Thus, the sampled type of the hospital θ_{it} is simply the set of donor/recipient pairs that were sampled at round t. At each round, the hospital decides between two strategies: in the truthful strategy, $Y_i(z,t) = 1$, the hospital reports $R_i(z,t) = \theta_{it}$. In the deviating strategy, $Y_i(z,t) = 0$ and the hospital performs all possible matches among its own pairs internally, and then reports the remainder $R_i(z,t) = d(\theta_{it})$. Given a pool of donor/patients θ_{it}, the function $d(.)$ is deterministic. Furthermore, as mentioned before, compatibility statistics are well-documented in the medical literature, and so the distributions of θ_{it} and $d(\theta_{it})$ (F and G respectively) are assumed known[5].

4.2 Simulation Setup

We perform simulation of two realistic stylized kidney exchange models that have been studied in literature. Mechanism $\mathcal{M}_0$ is the baseline mechanism and given a joint pool of hospital reports, computes a random maximum matching over all pairs. Mechanism $\mathcal{M}_1$ applies the *revelation principle* along with some more detailed allocations in well-defined subgroups of donor/patient

[5]For example, a pair with a donor with blood-type O is one that can possibly perform many exchanges, since O-donors can donate to all blood-types. Therefore, we expect more O-donors under distribution F than distribution G, since when deviating, hospitals are more likely to match these "good" pairs internally.

pairs[6]. Random graph theory has been leveraged to show that $\mathcal{M}_0$ is vulnerable to deviating hospitals while in $\mathcal{M}_1$ hospitals are better-off by being truthful (Toulis & Parkes, 2011).

As a ground-truth model of agent strategic behavior, we adopt a multi-armed bandit formulation. Specifically, we assume that hospitals try to maximize their utility (number of total pairs matched) by using the *uniform confidence bound* algorithm. This algorithm has been widely used in practice as the simplest and most effective model of dynamic strategic behavior with bounded regret (Auer et al., 2002). The algorithm used to generate the dataset is shown in Algorithm 3.

For $\forall$ hospital $i, \forall s \in \{0, 1\}$
  $u_{is} = 0$
  $n_{is} = 1$
For $t = 1, 2, \cdots T$
  For $i = 1, 2, \cdots N$
    $\theta_{it} \sim F$, sample the hospital's internal pairs
    $s^* = \arg\max_{s \in \{0,1\}} \left(u_{is} + \sqrt{\frac{2 \log t}{n_{is}}} \right)$
    $Y_i(z, t) = s^*$
    $n_{is^*} = n_{is^*} + 1$
    $R_i(Z_i, t) = \begin{cases} \theta_{it}, & \text{if } Y_i(z,t) = 1 \\ d(\theta_{it}), & \text{otherwise} \end{cases}$
  Causal effect estimand $\Delta(t) = h(\mathbf{Y}_1(t), \mathbf{Y}_0(t))$
Return $\Delta(t), \mathbf{R}_1^{obs}(t), \mathbf{R}_0^{obs}(t), \forall t$

Algorithm 3: Simulation model of dynamic hospital behavior in kidney exchanges. Variable n_{is} keeps track of how many times hospital i has chosen strategy s (1=truthful). Variable u_{is} keeps track of the average achieved utility of agent i by playing strategy s.

The output of the simulation of Algorithm 3 are the causal estimand values at different rounds t and the observed agent reports ($(N/2 \times 1)$ vectors of observed hospital reports for each t). Our inference goal for both the empirical and the game-theoretic methods will be to estimate the estimand values $\Delta(t)$ produced by the simulation, given the observed reports, for *two* specific rounds: $t_1 = 5$ will be round when the data are collected (observe agents' reports) and $t_2 = T = 100$ will be considered as the round where the system has reached equilibrium. Figure 1 shows 100 independent runs of the simulation with the multi-armed bandit dynamic and the confidence bands of the respective estimand values $\Delta(t)$ for every round t. Note that the causal effects to be estimated

[6]Briefly, pairs can be categorized as "under-demanded", "over-demanded", "reciprocal" and "self-demanded". The compatibility networks within these groups vary significantly. The "under-demanded" cannot be matched to each other and so the subgroup network is isolated. The "reciprocal" subgroup consists of two smaller groups that can be matched to each other and so the network is bipartite. The "self-demanded" is composed of four smaller groups that internally look like complete graphs. These nuisances can be leveraged for the design of better allocation mechanisms than myopic maximum matchings.

are time-dependent. For example, at round $t = 5$, the difference in incentives is around 0.1, specifically 0.5 for $\mathcal{M}_1$ and 0.4 for $\mathcal{M}_0$, which means that hospitals in $\mathcal{M}_1$ are 25% more likely to play the truthful strategy (t) compared to $\mathcal{M}_0$ at round 5 and under the multi-armed bandit dynamic. For larger t this value steadily increases until 0.44 and then stabilizes (based on additional experiments, we believe the seemingly linear trend at later time points, is an artifact of our simulation). The complication for causal inference is that unbiased estimates of incentives, taken at different timepoints, would be wildly different. This illustrates that estimation of mechanism effects needs to make the distinction between *short-term* and *long-term*. One key goal of this paper is to propose a methodology to estimate long-term effects by early experimental data. In our simulation, we assume that data are collected at $t = 5$ and we are interested in estimating the short term effect $\Delta(5)$ and long-term effect $\Delta(100)$.

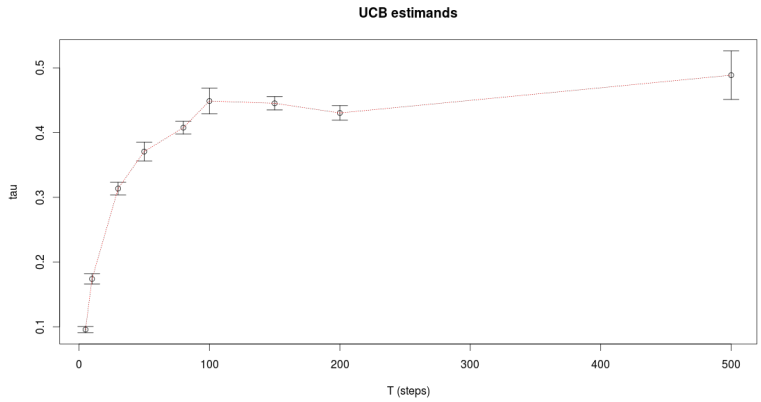

Figure 1: Causal estimand at different rounds t average over 100 independent runs. The estimand is time-dependent indicating that the system reaches equilibrium over time.

4.3 Estimation

We compare between the empirical method that uses uniform priors and our method which is using a game-theoretic prior. The former is straightforward and can be implemented by following Algorithm 1. The implementation of our method is a bit more involved as it first requires to compute the payoff functions $\Delta u(\cdot)$ as put forth in Equation (5) and under the exchangeability assumption (Assumption 3.2). In our simulation, this means that we have to calculate the payoff matrices for 9 cases. These can be obtained through simulations. The results over 10,000 mechanism simulations are shown in Table 2. For example, in $\mathcal{M}_0$ a truthful hospital will have an expected utility of 9.66 matches when there are 5 truthful hospitals out of $N = 8$. Thus if this hospital were to deviate, it would obtain an expected utility of 10.69 since the number of truthful hospital would decrease

Table 2: Payoff matrices of $\mathcal{M}_0$ and $\mathcal{M}_1$. Each cell in the matrix shows the expected utility (average number of matched pairs) given (i) a mechanism, (ii) number of truthful hospitals in the mechanism and (iii) an agent strategy. The table shows that in $\mathcal{M}_0$ it is better for hospitals to deviate and in $\mathcal{M}_1$ it is better to be truthful.

$\mathcal{M}_0$			$\mathcal{M}_1$		
expected utility			expected utility		
truthful	deviating	#truthful	truthful	deviating	#truthful
-	9.76	0	-	9.58	0
8.24	10.06	1	9.94	9.61	1
8.71	10.37	2	9.82	9.54	2
9.07	10.49	3	9.91	9.66	3
9.31	10.69	4	9.75	9.68	4
9.66	10.76	5	9.86	9.78	5
9.87	10.91	6	9.83	9.88	6
10.15	11.21	7	9.89	9.85	7
10.30	-	8	9.88	-	8

to 4. Table 2 shows that deviation is a dominant stratefy in $\mathcal{M}_0$ and truthful strategy is dominant strategy in $\mathcal{M}_1$.

Having obtained the payoff matrices, estimation through our model proceeds through the simple Gibbs sampling procedure described in Algorithm 3.

4.4 Results

We conduct two experiments and for each experiment we collect agent reports at $t = 5$ and wish to estimate short-term effects at $t = 5$ and long-term effects for $T = 100$. The ground truth estimands have simulated values $\Delta(5) = 0.1$ and $\Delta(100) = 0.44$ over 500 independent simulation runs.

In our first experiment, we work on a case in which agent reports are highly informative of the underlying agent behaviors. We refer to this case as "strong separability" since the posterior distributions $Y_i(z,t)|R_i(z,t)$ takes higher values (around 1) when the report is a truthful one and take small values (around 0) when the report is an untruthful one. In this cases the distributions look "separated" as in Figure 2a.

Figure 2b shows histograms of the estimates from the empirical method and the game-theoretic method. We can see that the former estimates $\Delta(5)$ at 0.12. Visual inspection of the estimates

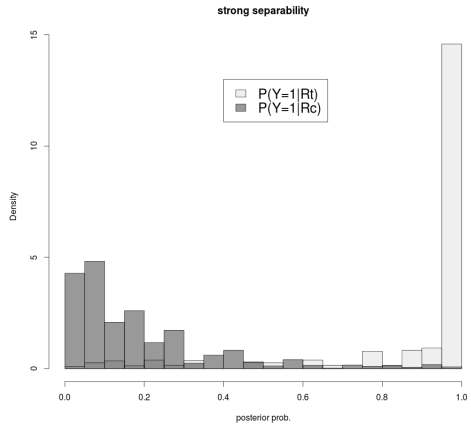

(a) Conditional distribution of agent strategy conditioned on truthful report (light gray) or untruthful report (dark gray).

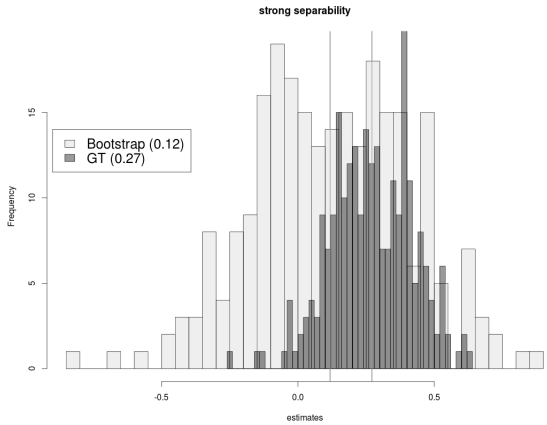

(b) Causal effects estimates from the empirical method (light gray) and the game-theoretic (GT) method (dark gray).

Figure 2: Experiment in which agent reports are informative of agent strategies (strong separability). Ground truth estimands have simulated values $\Delta(5) = 0.1$ and $\Delta(100) = 0.44$.

shows that the method performs well under the informative agent reports (strong separability). This is expected, since the likelihood gives plenty of information about the missing agent strategies. However, while the estimate (under the exchangeability assumption) is unbiased for the true effect $\Delta(5)$ at round $t = 5$, it is still biased for the long-term effect $\Delta(100)$. The game-theoretic method makes a compromise between the two estimands. The estimates are centered around 0.27 and they clearly biased for the effect at $t = 5$. However, they are also shrinked towards the long-term effect $\Delta(100)$ which indicates that the payoff matrix is able to capture the dynamic evolution of the system to a certain extent.

In our second experiment, we work on a case in which agent reports are not informative about the underlying agent behaviors. We refer to this case as "weak separability" since the posterior distributions $Y_i(z,t)|R_i(z,t)$ takes values around 0.5 for both truthful and untruthful reports so that there is practically no information about strategies given observed agents' reports. The respective distributions are shown in Figure 3a.

Figure 3b shows histograms of the estimates from the empirical method and the game-theoretic method for the weak separability case. We can see that the former estimates $\Delta(5)$ at 0.02. In that case, we observe a breakdown of the empirical method since the estimates are centered around zero. This is because the reports are not informative (almost random guesses) about the strategies and so the empirical method is using only the information on the prior to make inference about incentives. However, since the empirical method is assuming a uniform prior, the overall procedure will deduce no difference in incentives. In constrast, the game-theoretic method is actually giving

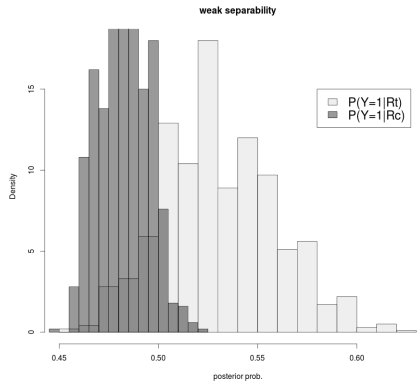

(a) Conditional distribution of agent strategy conditioned on truthful report (light gray) or untruthful report (dark gray).

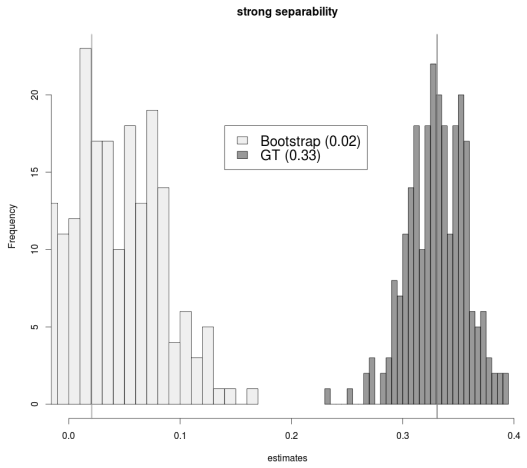

(b) Causal effects estimates from the empirical method (light gray) and the game-theoretic (GT) method (dark gray).

Figure 3: Experiment in which agent reports are informative of agent strategies (strong separability). Ground truth estimands have simulated values $\Delta(5) = 0.1$ and $\Delta(100) = 0.44$.

higher estimates for the overall difference in incentives since it is based only on the equilibrium model of the prior. Thus, the overall estimates are even higher than before (average around 0.33) as the equilibrium model shrinks the estimates towards long-term effects.

5 Discussion

The evaluation of mechanisms is critical in numerous socioeconomic problems. However, this is technically challenging because multi-agent systems are dynamic by nature and estimation should be performed with respect to an equilbrium state of the system. Furthermore, in estimating effects of incentives, there are additional challenges as agent strategies, which are the main potential outcomes of interest, are typically never observed. For the former, we use agent reports and further distributional assumptions to obtain likelihoods of strategies given observed reports. For the latter, we propose a prior on agent strategies that is based on a quantal response equilibrium model. In a simulated study, this was shown to shrink towards long-term effects, thus offering improved inference over methods that don't consider such priors.

There are multiple ways that this work could be further improved. First, the empirical method we described is by now means the only way that could be used to perform causal inference. A fully Bayesian model would also be a good choice. However, this work hints that any model that is

ignoring the game-theoretic aspect of the potential outcomes (agent behaviors) will be inadequate to capture the time dependence of the causal estimand. Second, the choice of the β parameter in the game-theoretic prior is crucial and how it was set, was not sufficiently justified. In practice, β was set based on a heuristic calculation and experimental results. Future work would be benefited by a more principled way to set such hyperparameters. Third, we offered limited discussion of our assumptions (best-response and exchangeability). Several violations of these assumptions yield more realistic and particularly interesting situations. For example, in case of substitution effects i.e., cases where agents can switch between mechanisms (e.g. assume that a mechanism is a mode of transportation), it is no longer possible to model the two potential outcome vectors independently. Last but not least, agent interactions (e.g. information sharing, communication, collusion) will require more sophisticated models.

References

ASHLAGI, I., FISCHER, F., KASH, I. & PROCACCIA, A. D. (2010). Mix and match. In *Proceedings of the 11th ACM conference on Electronic commerce*. ACM.

ASHLAGI, I. & ROTH, A. (2011). Individual rationality and participation in large scale, multi-hospital kidney exchange. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM.

ATHEY, S. & NEKIPELOV, D. (2010). A structural model of sponsored search advertising auctions. In *Sixth Ad Auctions Workshop*.

AUER, P., CESA-BIANCHI, N. & FISCHER, P. (2002). Finite-time analysis of the multiarmed bandit problem. *Machine learning* **47**, 235–256.

BOTTOU, L., PETERS, J., QUIÑONERO-CANDELA, J., CHARLES, D. X., CHICKERING, D. M., PORTUGUALY, E., RAY, D., SIMARD, P. & SNELSON, E. (2012). Couterfactual reasoning and learning systems. *arXiv preprint arXiv:1209.2355* .

CLARKE, E. H. (1971). Multipart pricing of public goods. *Public choice* **11**, 17–33.

GOEREE, J. K., HOLT, C. A. & PALFREY, T. R. (2003). Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior* **45**, 97–113.

GROVES, T. (1973). Incentives in teams. *Econometrica: Journal of the Econometric Society* , 617–631.

HECKMAN, J. J. & VYTLACIL, E. (2005). Structural equations, treatment effects, and econometric policy evaluation1. *Econometrica* **73**, 669–738.

LUBIN, B. & PARKES, D. C. (2009). Quantifying the strategyproofness of mechanisms via metrics on payoff distributions. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*. AUAI Press.

MCKELVEY, R. D. & PALFREY, T. R. (1995). Quantal response equilibria for normal form games. *Games and economic behavior* **10**, 6–38.

OSTROVSKY, M. & SCHWARZ, M. (2011). Reserve prices in internet advertising auctions: A field experiment. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM.

ROTH, A. E., SÖNMEZ, T. & ÜNVER, M. U. (2004). Kidney exchange. *The Quarterly Journal of Economics* **119**, 457–488.

RUBIN, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology* **66**, 688.

RUBIN, D. B. (1978). Bayesian inference for causal effects: The role of randomization. *The Annals of Statistics* , 34–58.

TOULIS, P. & PARKES, D. C. (2011). A random graph model of kidney exchanges: efficiency, individual-rationality and incentives. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM.

VICKREY, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance* **16**, 8–37.

PANOS TOULIS, Department of Statistics, Harvard University

E-mail address: ptoulis@fas.harvard.edu

DAVID C. PARKES, SEAS, Harvard University

E-mail address: parkes@eecs.harvard.edu